Multi-study Factor Analysis

Roberta De Vito*[1], Ruggero Bellio[2], Lorenzo Trippa[3,4], and Giovanni Parmigiani[3,4]

[1]*Department of Computer Science, Princeton University, Princeton, NJ, USA*
[2]*Department of Economics and Statistics, University of Udine, Udine, Italy*
[3]*Department of Biostatistics and Computational Biology, Dana-Farber Cancer Institute, Boston, MA, USA*
[4]*Department of Biostatistics, Harvard T. H. Chan School of Public Health, Boston, MA, USA*

arXiv:1611.06350v2 [stat.AP] 4 Jan 2017

January 5, 2017

Abstract

We introduce a novel class of factor analysis methodologies for the joint analysis of multiple studies. The goal is to separately identify and estimate 1) common factors shared across multiple studies, and 2) study-specific factors. We develop a fast Expectation Conditional-Maximization algorithm for parameter estimates and we provide a procedure for choosing the common and specific factor. We present simulations evaluating the performance of the method and we illustrate it by applying it to gene expression data in ovarian cancer. In both cases, we clarify the benefits of a joint analysis compared to the standard factor analysis. We hope to have provided a valuable tool to accelerate the pace at which we can combine unsupervised analysis across multiple studies, and understand the cross-study reproducibility of signal in multivariate data.

1 Introduction

High-dimensional data are the norm in current statistical research. Building systematic knowledge from these data is a cumulative process, which requires analyses that integrate multiple sources, studies, and data-collection technologies. A common goal in high-dimensional data analysis is the unsupervised identification of latent components or factors. When considering multiple studies, a fundamental challenge is learning common features shared among studies while isolating the variation specific to each study. Two important statistical questions remain largely unanswered in this context: i) To what extent is signal shared across studies? ii) How can shared signal be extracted? In this paper we develop a multi-study factor analysis methodology to address these questions.

*rvito@princeton.edu

Joint factor analysis of multiple studies is common in several areas of science. For example Scaramella et al. (2002) study adolescent delinquent behavior in two independent samples to identify shared patterns. In nutritional epidemiology Edefonti et al. (2012) used a factor analysis (FA) across five different studies to determine common dietary patterns and their relation with head and neck cancer. Andreasen et al. (2005) studied remission in schizophrenia, showing that replicable results are found across FA leading to similar components. Wang et al. (2011) used FA to obtain a unified gene expression measurement from multiple types of measurements on the same samples.

Our own motivating example derives from the analysis of gene expression data (Irizarry et al., 2003, Shi et al., 2006, Kerr, 2007). In gene expression analysis, as well as in much of high-throughput biology analyses on human populations, variation can arise from the intrinsic biological heterogeneity of the populations being studied, or from technological differences in data acquisition. In turn both these types of variation can be shared across studies or not. As noted by Garrett-Mayer et al. (2007), the fact that the determinants of both natural and technological variation differ across studies implies that study-specific effects occur in most datasets. Both common and study-specific effects can be strong, and both need to be identified and studied. Our interest in this issue is a natural development of our previous work on unsupervised identification of integrative correlation (Parmigiani et al., 2004, Garrett-Mayer et al., 2007, Cope et al., 2014), and multi-study supervised analyses including cross-study differential expression (Scharpf et al., 2009), multi-study gene set analysis (Tyekucheva et al., 2011) comparative meta-analysis (Riester et al., 2014, Waldron et al., 2014), and cross study validation (Bernau et al., 2014).

In high throughput biology, as well as in a number of other areas of application, the ability to separately estimate common and study-specific factors can contribute significantly to two important questions: the cross-study validation question of whether factors are found repeatedly across multiple studies; and the meta-analytic question of more efficiently estimating the factors that are indeed common. With regard to interpretation, the shared signal is more likely to capture genuine biological information, while the study-specific signal can point to either artifactual or biological sources of variation. Thus, model both shared and unshared factors may enable a more reliable identification of artifacts, facilitate more efficient experimental designs, and inform further technological advances.

In this article we propose a dimension-reduction approach that allows for joint analysis of multiple studies, achieving the goal of capturing common factors. Specifically, we define a generalized version of FA, able to handle multiple studies simultaneously. Our model, termed Multi-study Factor Analysis (MSFA), learns the common features shared among studies, and identifies the unique variation present in each study.

While unsupervised multi-study analysis is not an adequately studied field, our work draws from existing foundations from related problems. In the social science literature, there is extensive methodology to find factor structures shared among different groups, forming the body of *multigroup factor analysis* methods (see, among many others, Thurstone (1931), Jöreskog (1971), Meredith (1993)). Such methods are mainly focused on investigating measurement invariance among different groups, that typically results in testing whether the data support the hypothesis of a common loading matrix across groups. A notable special case is given by partial measurement invariance (see for example Byrne et al., 1989), which inspired our mathematical for-

Table 1: The four data sets considered in our illustration and their characteristics.

Study	Samples	Platform	Late Stage (%)	Reference
GSE9891	285	Affy U133Plus 2.0	85	Tothill et al. (2008)
GSE20565	140	Affy U133Plus 2.0	48	Meyniel et al. (2010)
GSE26712	195	Affy U133a	96	Bonome et al. (2008)
TCGA	578	Affy HT U133a	90	Network et al. (2011)

mulation. In our MSFA we have extended the scope to detection of both study-specific factors and factors that are identical across multiple studies. Our MSFA has also an exploratory nature, different from the confirmative approach under which measurement invariance is usually investigated in the social sciences.

The MSFA is applied here in the context of genomic data. In previous literature FA has been used in several gene expression studies. Among many others,Blum et al. (2010) adopted FA for multiple testing developed by Friguet et al. (2009) to characterize patterns of heterogeneity in gene expression data sets, and Runcie and Mukherjee (2013) used a sparse factor model by Bhattacharya and Dunson (2011) to capture subsets of important biological factors that control the variation in high-dimensional phenotypes. The most notable difference between the approach proposed here and its predecessors is the focus on distinguishing shared from study-specific factors.

The plan of the paper is as follows. Section 2 introduces the MSFA, and describes the estimation of model parameters based on maximum likelihood, implemented via an ECM algorithm. Section 3 presents simulation studies, providing numerical evidence on the performances of the proposed estimation methods. Next it investigates determining the dimension of the latent factor, solved by casting it within the framework of model selection. Section 4 illustrates the application of the methodology to study of the Immune System pathway in ovarian cancer. Section 5 contains the final discussions.

2 Methods

2.1 Gene Expression Illustration

For clarity, we present the method weaving through an illustrative analysis of gene expression data from ovarian cancer samples, the genomic application that motivated its development. We emphasize that the model can be applied to a variety of other situations where both similarities and differences exist in multiple high dimensional data sets. In biology alone, the approach proposed in this work can be applied to a large and diverse collection of high-throughput studies ranging from gene expression via microarrays or RNA-seq, to genome-wide association studies (GWAS).

We use four publicly available ovarian cancer (OC) microarray gene expression studies (Ganzfried et al., 2013), from the `curatedOvarianData` package in `Bioconductor`. Table 1 provides an overview of the studies with corresponding references, information about the specific microarray platform used, and the proportion of patients diagnosed with Stage III or Stage IV OC.

The total sample size is $n = 1198$. In our illustration, we focus on genes assigned to the Immune System pathway. The Immune System is of particular interest, as knowledge gained can drive the development of targeted therapies and tumor marker-

based diagnostic tests (Méhes et al., 2001). For each study, we only consider the Immune System genes measured in all four studies.

2.2 The multi-study factor analysis (MSFA) model

The methodology proposed here has two main goals. First, we combine multiple studies to identify common factors that are consistently observed across the studies. Second, we explicitly model and identify additional components of variability specific to single studies, through study-specific factors. The latter aims to identify variation that lacks cross-study reproducibility, and separate it from variation that does.

We consider S studies, each with the same P variables. Generic study s has n_s subjects and, for each subject, P-dimensional centered data vector $\mathbf{x}_{is}$ with $i = 1, \ldots, n_s$. We begin by describing the case where a standard FA is carried out separately in each study. The observed variables in study s are decomposed into J_s factors related to that study's specific sources of variation. Factor loadings relate linearly the observed variables to the latent factors. In particular, let $\mathbf{l}_{is}$, $i = 1, \ldots, n_s$ be the values of the *study-specific* factors in individual i of study s and $\mathbf{\Lambda}_s$, $s = 1, \ldots, S$ be the $P \times J_s$ corresponding factor loading matrix. FA assumes that the vector $\mathbf{x}_{is}$ is decomposed as

$$\mathbf{x}_{is} = \mathbf{\Lambda}_s \mathbf{l}_{is} + \mathbf{e}_{is} \quad i = 1, \ldots, n_s \,, \tag{1}$$

where $\mathbf{e}_{is}$ is a Gaussian error term with covariance matrix $\mathbf{\Psi}_s = \text{diag}(\psi_{1s}, \ldots, \psi_{ps})$ (Jöreskog, 1967a,b, Jöreskog and Goldberger, 1972). Equivalently, FA aims at explaining the dependence structure among high-dimensional observations by decomposing the $P \times P$ covariance matrix $\mathbf{\Sigma}_s$ as

$$\mathbf{\Sigma}_s = \mathbf{\Lambda}_s \mathbf{\Lambda}_s^\top + \mathbf{\Psi}_s \,. \tag{2}$$

In our illustration, we performed standard FA, as in equation (1), separately for the first two studies in Table 1: Tothill et al. (2008) and Meyniel et al. (2010). We estimate parameters by maximum likelihood, implemented via the EM algorithm using the identifying constraint that the loading matrix is lower triangular (Geweke and Zhou, 1996, Lopes and West, 2004, Carvalho et al., 2008).

Figure 1: Left: Heatmap of the factor loadings estimated by performing a separate factor analysis in studies GSE9891 and GSE20565 from Table 1. Right: Graphical representation of the absolute value of correlations between pairs of study-specific factor loadings. Correlations smaller than .2 are not shown.

The heatmap in Figure 1 shows the estimated factor loadings, $\mathbf{\Lambda}_1 = \{\lambda_{11}, \lambda_{21}, \lambda_{31}, \lambda_{41}, \lambda_{51}, \lambda_{61}\}$ for the GSE9891 study and $\mathbf{\Lambda}_2 = \{\lambda_{12}, \lambda_{22}, \lambda_{32}, \lambda_{42}, \lambda_{52}, \lambda_{62}, \lambda_{72}\}$ for the

GSE20565 study, obtained by performing a separated FA in each study. Each column λ_{is} is thus the i^{th} loading vector of the s^{th} study.

The left panel of Figure 1 suggests that some loading vectors may reveal common patterns across studies. This point is further explored in the right panel, obtained from the cross-study pairwise correlation of the loading vectors for the two studies. In these correlation the same variables are considered in each study, and their order is preserved. In the right panel of Figure 1 darker lines denote larger correlations in absolute value. Three of the loading vectors of the GSE9891 study are strongly correlated with four factors in the GSE20565 study.

Highly correlated pairs of loading vectors are more likely to represent common factors. On the other hand, some loading vectors of GSE9891 (e.g. λ_{41}) exhibit low correlation with all loading vectors of GSE20565. These loadings are likely to result from feature unique to this study. From a statistical perspective, it is important to focus on the common features shared among the studies, and on their potential connection with biological signal. A second important goal is the interpretation of study-specific sources of variations that are only present in a single dataset.

The MSFA model proposed here is designed to analyze multiple studies jointly, replacing the heuristic interpretation above with a principled statistical approach. It explicitly models common biological features shared among the studies, as well as unique variation present in each study. Specifically, the observed variables in study s are decomposed into K factors shared with the other studies, and J_s additional factors reflecting its unique sources of variation, for a total of $K + J_s$ factors. Let $\mathbf{f}_{is}$ be the *common* factor vector in subject i of study s, and $\mathbf{\Phi}$ be the $P \times K$ common factors loading matrix. Moreover, let $\mathbf{l}_{is}$ be the *study-specific* factor and $\mathbf{\Lambda}_s$ be the $P \times J_s$ specific factors loading matrix. MSFA assumes that the P-dimensional centered response vector $\mathbf{x}_{is}$ can be written as

$$\mathbf{x}_{is} = \mathbf{\Phi}\mathbf{f}_{is} + \mathbf{\Lambda}_s\mathbf{l}_{is} + \mathbf{e}_{is}\,, \qquad i = 1, \ldots, n_s \qquad s = 1, \ldots, S. \tag{3}$$

where the $p \times 1$ random error vector $\mathbf{e}_{is}$ has a multivariate normal distribution with mean vector $\mathbf{0}$ and covariance matrix $\mathbf{\Psi}_s$, with $\mathbf{\Psi}_s = \mathrm{diag}(\psi^2_{s_1}, \ldots, \psi^2_{s_p})$.

We also assume that the marginal distribution of $\mathbf{l}_{is}$ is multivariate normal with mean vector $\mathbf{0}$ and covariance matrix $\mathbf{I}_{J_s}$, and the marginal distribution of $\mathbf{f}_{is}$ is multivariate normal with mean vector $\mathbf{0}$ and covariance matrix $\mathbf{I}_k$, where $\mathbf{I}$ denotes the identity matrix. The left panel of Figure 2 represents our model (3) graphically.

As a result of our assumptions, the marginal distribution of $\mathbf{x}_{is}$ is multivariate normal with mean vector $\mathbf{0}$ and covariance matrix $\mathbf{\Sigma}_s = \mathbf{\Phi}\mathbf{\Phi}^\top + \mathbf{\Lambda}_s\mathbf{\Lambda}_s^\top + \mathbf{\Psi}_s$, with the three terms reflecting the variance of the common factors, the variance of the study-specific factors, and the variance of the error.

This model can be applied to many settings when the aim is to isolate commonalities from differences across groups, population or studies. In the gene expression application the focus is on estimating the biological signal shared among the studies, while removing study-specific features less likely to be reproducible across populations, and potentially arising from technological issues. This is only one of the applications of the MSFA model. Elsewhere the goal may be to capture some study-specific features of interest while removing common factors shared among the studies. Yet other applications may focus on studying both common and specific factors.

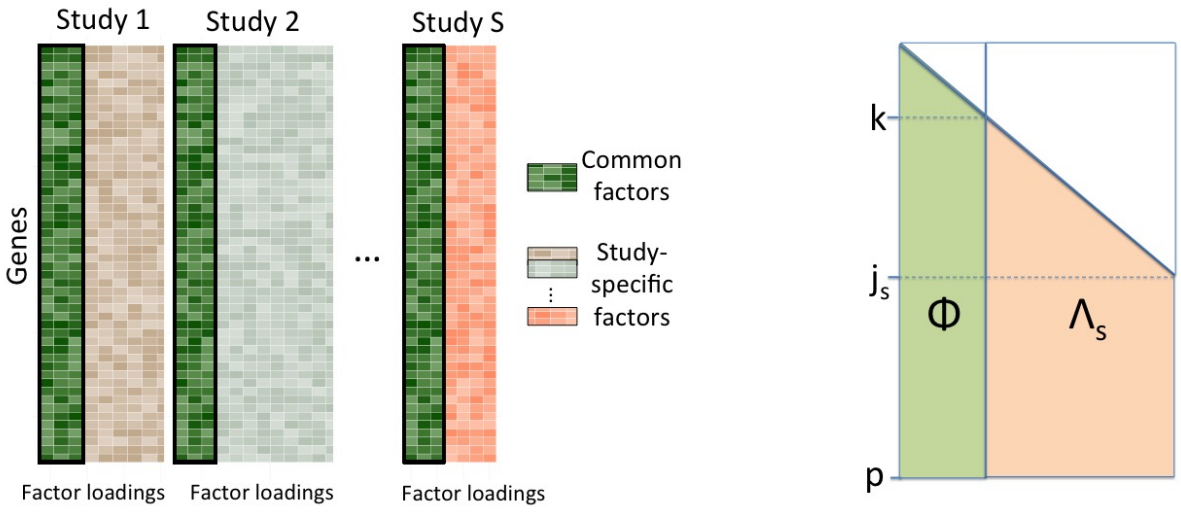

Figure 2: Left: schematic of the decomposition of the latent factors matrices $\mathbf{\Omega}_s$ into common (in green) and study-specific components, without the constraints. Right: schematic of the the identifying constraints on each $\mathbf{\Omega}_s$.

2.3 Identifiability

To specify an identifiable model, the MSFA model must be further constrained to avoid orthogonal rotation indeterminacy, similarly to the classic FA model. Lack of identifiability can be seen by considering the model for a subject in the s^{th} study. Let $\mathbf{\Omega}_s = [\mathbf{\Phi}, \mathbf{\Lambda}_s]$ be the $P \times (K + J_s)$ loading matrix for the s^{th} study. If we define $\mathbf{\Omega}_s^* = \mathbf{\Omega}_s \mathbf{Q}_s$, where $\mathbf{Q}_s$ is a square orthogonal matrix with $(K + J_s)$ rows, it readily follows that $\mathbf{\Omega}_s^*(\mathbf{\Omega}_s^*)^\top = \mathbf{\Omega}_s \mathbf{Q}_s \mathbf{Q}_s^\top \mathbf{\Omega}_s^\top = \mathbf{\Omega}_s \mathbf{\Omega}_s^\top$, so that

$$\mathbf{\Sigma}_s = \mathbf{\Omega}_s^*(\mathbf{\Omega}_s^*)^\top + \mathbf{\Psi}_s = \mathbf{\Omega}_s \mathbf{\Omega}_s^\top + \mathbf{\Psi}_s \,,$$

and $\mathbf{\Sigma}_s$ is not uniquely identified.

The standard factor model (1) identifies the parameters by imposing constraints on the matrix of factor loadings. One possibility often used in practice is to take $\mathbf{\Lambda}_s$ in (1) as a lower triangular matrix (Geweke and Zhou, 1996, Lopes and West, 2004). Here we adapt this approach to the MSFA model, and specify $\mathbf{\Omega}_s = [\mathbf{\Phi}, \mathbf{\Lambda}_s]$ to be lower triangular, as illustrated in the right panel of Figure 2. Importantly, the matrices $\mathbf{\Phi}$ and $\mathbf{\Lambda}_s$ are not interchangeable

As in the standard FA model, assuming a lower triangular form for $\mathbf{\Omega}_s$ resolves the orthogonal rotation indeterminacy, for the same reason explained in Geweke and Zhou (1996, pp. 565-566). There remain the issue that we can simultaneously change the sign to all the elements of the loading matrices, and to all the latent factors without changing the model. This could be fixed by constraining the sign of a subset of loadings, but for maximum likelihood estimation of the model parameters this issue is largely inconsequential, as will be discussed later.

We denote by $\mathbf{C}_{x_s x_s}$ the sample covariance matrix for the s^{th} study. The total number of elements in $\mathbf{C}_{x_s x_s}$ must be no greater than the number of free parameters in the covariance matrix $\mathbf{\Sigma}_s$, implying that the following set of S conditions must hold

$$P(K + J_s) + P - \frac{(K + J_s)(K + J_s - 1)}{2} \leq \frac{1}{2} P(P + 1) \,, \qquad s = 1, \ldots, S \,.$$

2.4 Parameter estimation

The parameters to be estimated in the MSFA are $\boldsymbol{\theta} = (\boldsymbol{\Phi}, \boldsymbol{\Lambda}_s, \boldsymbol{\Psi}_s)$. For notational simplicity in both (1) and (3) we assume that the observed variables in each study have been centered at the sample means. The log-likelihood function corresponding to the MSFA assumptions is

$$\ell(\boldsymbol{\theta}) \;=\; \log \prod_{s=1}^{S} \prod_{i=1}^{n_s} p(\mathbf{x}_{is}|\boldsymbol{\theta}) = \sum_{s=1}^{S} \left\{ -\frac{n_s}{2} \log |\boldsymbol{\Sigma}_s| - \frac{n_s}{2} \mathrm{tr}(\boldsymbol{\Sigma}_s^{-1} \mathbf{C}_{x_s x_s}) \right\} .$$

In the following, we show how to compute the Maximum Likelihood Estimate (MLE) for the MSFA model by the Expectation Conditional Maximization (ECM) algorithm.

2.4.1 MLE computation using the ECM algorithm

The Expectation-Maximization (EM) algorithm (Dempster et al., 1977) is now a standard technique for obtaining the MLE in factor analysis models (Rubin and Thayer, 1982). If $\mathbf{f}_{is}$ and $\mathbf{l}_{is}$ were observed, the log-likelihood would be

$$l_c(\boldsymbol{\theta}) = \log \prod_{s=1}^{S} \prod_{i=1}^{n_s} p(\mathbf{x}_{is}, \mathbf{f}_{is}, \mathbf{l}_{is}|\boldsymbol{\theta}) \,. \tag{4}$$

There are two steps in each iteration t of the EM algorithm. First, in the E step, we find the expectation of $l_c(\boldsymbol{\theta})$ by integrating over the latent variables, $\mathbf{f}_{is}$ and $\mathbf{l}_{is}$ conditional distributions, with the parameter $\boldsymbol{\theta}$ held fixed at the value $\boldsymbol{\theta}_{t-1}$ of the previous iteration. The second step, the M-step, maximizes the expected log-likelihood, and yields the next value $\boldsymbol{\theta}_t$ of the parameter.

The ECM algorithm by Meng and Rubin (1993) replaces the maximization step of the EM algorithm with a set of sequential maximization steps, considering each parameter (or group of parameters) in turn, while keeping the other fixed. This reduces a complex maximization problem into several simpler ones. The relevant convergence properties of the EM algorithm are retained by the ECM algorithm (McLachlan and Krishnan, 2007, §1.7).

Finding the expectation of $l_c(\boldsymbol{\theta})$ given $\mathbf{x}_{is}$ for each $s = 1, \ldots, S$ and $\boldsymbol{\theta}_{t-1}$ requires finding the conditional expectation of the following quantities

$$\begin{array}{lll} \mathbf{C}_{x_s x_s} = \dfrac{\sum_{i=1}^{n_s} \mathbf{x}_{is}\mathbf{x}_{is}^\top}{n_s}, & \mathbf{C}_{ff} = \dfrac{\sum_{i=1}^{n_s} \mathbf{f}_{is}\mathbf{f}_{is}^\top}{n_s}, & \mathbf{C}_{l_s l_s} = \dfrac{\sum_{i=1}^{n_s} \mathbf{l}_{is}\mathbf{l}_{is}^\top}{n_s}, \\ \mathbf{C}_{x_s f} = \dfrac{\sum_{i=1}^{n_s} \mathbf{x}_{is}\mathbf{f}_{is}^\top}{n_s}, & \mathbf{C}_{x_s l_s} = \dfrac{\sum_{i=1}^{n_s} \mathbf{x}_{is}\mathbf{l}_{is}^\top}{n_s}, & \mathbf{C}_{f l_s} = \dfrac{\sum_{i=1}^{n_s} \mathbf{f}_{is}\mathbf{l}_{is}^\top}{n_s}. \end{array}$$

given the parameters and the observed data. These are:

$$\begin{array}{ll} \mathbf{T}_{x_s x_s} = \mathrm{E}\left[\mathbf{C}_{\mathrm{x_s x_s}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \mathbf{C}_{\mathrm{x_s x_s}}, & \mathbf{T}_{ff} = \mathrm{E}\left[\mathbf{C}_{\mathrm{ff}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \boldsymbol{\delta}\mathbf{C}_{\mathrm{x_s x_s}}\boldsymbol{\delta}^\top + \boldsymbol{\Delta}, \\ \mathbf{T}_{x_s l_s} = \mathrm{E}\left[\mathbf{C}_{\mathrm{x_s l_s}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \mathbf{C}_{\mathrm{x_s x_s}}\boldsymbol{\delta}_{\mathrm{s}}^\top, & \mathbf{T}_{l_s l_s} = \mathrm{E}\left[\mathbf{C}_{\mathrm{l_s l_s}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \boldsymbol{\delta}_{\mathrm{s}}\mathbf{C}_{\mathrm{x_s x_s}}\boldsymbol{\delta}_{\mathrm{s}}^\top + \boldsymbol{\Delta}_{\mathrm{s}}, \\ \mathbf{T}_{x_s f} = \mathrm{E}\left[\mathbf{C}_{\mathrm{x_s f}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \mathbf{C}_{\mathrm{x_s x_s}}\boldsymbol{\delta}^\top, & \mathbf{T}_{f l_s} = \mathrm{E}\left[\mathbf{C}_{\mathrm{f l_s}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}\right] = \boldsymbol{\delta}\mathbf{C}_{\mathrm{x_s x_s}}\boldsymbol{\delta}_{\mathrm{s}}^\top + \boldsymbol{\Delta}_{\mathrm{fl}}. \end{array}$$

where $\boldsymbol{\Delta}_{fl} = \mathrm{Cov}[\mathbf{l}_{\mathrm{is}}, \mathbf{f}_{\mathrm{is}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}]$ and

$$\begin{array}{ll} \boldsymbol{\delta} = \boldsymbol{\Phi}^\top \boldsymbol{\Sigma}_s^{-1}, & \boldsymbol{\Delta} = \mathrm{Var}[\mathbf{f}_{\mathrm{is}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}] = \mathbf{I}_{\mathrm{k}} - \boldsymbol{\Phi}^\top \boldsymbol{\Sigma}_{\mathrm{s}}^{-1} \boldsymbol{\Phi}, \\ \boldsymbol{\delta}_s = \boldsymbol{\Lambda}_s^\top \boldsymbol{\Sigma}_s^{-1}, & \boldsymbol{\Delta}_s = \mathrm{Var}[\mathbf{l}_{\mathrm{is}}|\mathbf{x}_{\mathrm{is}}, \boldsymbol{\theta}_{\mathrm{t-1}}] = \mathbf{I}_{\mathrm{j_s}} - \boldsymbol{\Lambda}_{\mathrm{s}}^\top \boldsymbol{\Sigma}_{\mathrm{s}}^{-1} \boldsymbol{\Lambda}_{\mathrm{s}}. \end{array}$$

The CM update is then carried out in three steps, each updating a set of parameter while keeping the others at their most recent values.

CM1 updates of the error covariance matrix $\mathbf{\Psi}_s$:

$$\mathbf{\Psi}_s^{new} = \text{diag}\left(\mathbf{C}_{\mathrm{x_s x_s}} + \mathbf{\Phi}\mathbf{T}_{\mathrm{ff}}\mathbf{\Phi}^\top + \mathbf{\Lambda}_{\mathrm{s}}\mathbf{T}_{\mathrm{l_s l_s}}\mathbf{\Lambda}_{\mathrm{s}}^\top - 2\mathbf{T}_{\mathrm{x_s f}}\mathbf{\Lambda}_{\mathrm{s}}^\top - 2\mathbf{T}_{\mathrm{x_s f_s}}\mathbf{\Lambda}_{\mathrm{s}}^\top + 2\mathbf{\Phi}\mathbf{T}_{\mathrm{fl_s}}\mathbf{\Lambda}_{\mathrm{s}}^\top\right).$$

CM2 updates $\mathbf{\Phi}$:

$$\text{vec}(\mathbf{\Phi}^{new}) = \sum_{s=1}^{S}\left(\mathbf{T}_{ff}^\top \otimes n_s \mathbf{\Psi}_s^{new^{-1}}\right)\text{vec}\left(n_s\mathbf{\Psi}_s^{new^{-1}}\mathbf{T}_{x_s f} - n_s\mathbf{\Psi}_s^{new^{-1}}\mathbf{\Lambda}_s\mathbf{T}_{fl_s}^\top\right).$$

Here $\otimes$ is the Kronecker product, vec is the vec operator and the linear equation is solved by the Lyapunov Equation (e.g. Van Loan, 2000).

CM3 updates $\mathbf{\Lambda}_s$:

$$\mathbf{\Lambda}_s^{new} \;=\; \left(\mathbf{T}_{x_s l_s} - \mathbf{\Phi}^{new}\mathbf{T}_{f l_s}\right)\left(\mathbf{T}_{l_s l_s}\right)^{-1}.$$

To assess convergence, we used the Aitken acceleration-based stopping criterion (McLachlan and Krishnan, 2007, §4.9). For simplicity of notation, let $l_c(\boldsymbol{\theta}^{(t+1)}) = l^{(t+1)}$ be the complete log-likelihood obtained at $(t+1)$ iteration and $l_A^{(t+1)} = l^{(t)} + \frac{1}{1-c^{(t)}}(l^{(t+1)} - l^{(t)})$, where $c^{(t)} = (l^{(t+1)} - l^{(t)})/(l^{(t)} - l^{(t-1)})$. We stop the ECM if $|l_A^{(t+1)} - l_A^{(t)}|$ becomes smaller than a specified threshold.

2.4.2 Dimension Selection

Selecting the dimension of the model can be challenging. In our simulations and applications we use the following two-step procedure. First we determine the total latent dimension $T_s = K + J_s$ for each of the S studies using the techniques for the standard FA, such as Horn's parallel analysis (Horn, 1965), Cattell's scree test (Cattell, 1966) or the use of indexes, such as the RMSEA (Steiger and Lind, 1980). Next, we use model selection techniques on the overall MSFA model to select the number K of latent factors sharing a common loading matrix $\mathbf{\Phi}$. The dimensions J_s are then obtained residually as $T_s - K$, $s = 1, \ldots, S$.

3 Simulation studies

We present next our simulation experiments designed to evaluate the ECM algorithm performance in estimating the MSFA model parameters, and our strategy for selecting the dimensionality of the latent factors.

We design our simulations to closely mimick the data of Table 1. We consider $S = 4$ studies, with dimension of the latent factors as in Table 2. We consider the same $p = 100$ variables in each study. We study three simulation scenarios. In Scenario 1 there are no common factors, i.e. $K = 0$, in Scenario 2 we set $K = 1$, and in Scenario 3 we set $K = 3$. For realism, in each case, the data are generated by parameter values akin to those estimated with the data.

Table 2: Settings for the simulation studies.

s	n_s	$K + J_s$
1	285	6
2	140	7
3	195	11
4	578	10

3.1 Parameter estimation via the ECM algorithm

We first analyze the performances of the ECM algorithm for a given selection of K and J_s, $s = 1, \ldots, S$. We compare the results obtained with the ECM algorithm to those obtained with the box-constrained Limited-memory Broyden-Fletcher-Goldfarb-Shannon (L-BFGS) method (e.g. Byrd et al., 1995). This method is used for benchmarking the ECM algorithm in its standard form, without any particular effort to tailor the optimization to the model at hand.

Irrespective of the optimization method adopted, the choice of the starting point is crucial for achieving good performances. We used following strategy, for given factor dimensions K and J_s, $s = 1, \ldots, S$.

1. A single data set is created by stacking the data of the four studies by row, obtaining a single data set with $n = n_1+n_2+n_3+n_4 = 578+285+195+140 = 1198$ observations and $p = 100$ variables.

2. A Principal Components Analysis (PCA) is performed on the data set obtained at the first step. The first K principal components are taken as the starting point of the common factor loadings.

3. Transform the data to remove the variability represented by the first K principal components.

4. Fit a standard FA model to the transformed data, for each study separately. The factor loadings and uniquenesses matrix obtained from the standard FA are used as the initial values for $\mathbf{\Lambda}_s$ and $\mathbf{\Psi}_s$.

Figure 3 considers three illustrative data sets, and illustrates the ascent of log-likelihood function as iterations progress, after starting the two algorithms from the same point. These examples highlight the much faster convergence for the ECM algorithm and are represented Ultimately, convergence is attained for both methods, but for the L-BFGS the required number of iterations is much higher, around 20,000. Moreover, with the L-BFGS, we began to experience severe convergence problems when used for larger number of variables.

Adachi (2012) reported that the EM algorithm in FA models always gives proper solutions when the sample covariance and initial parameter matrices are proper. Empirically, our experience has been that the same result is true of ECM for the MSFA model, at least in the simulated data sets we studied.

Next we consider the Likelihood Ratio Test (LRT) for choosing between $K = 0$ and $K = 1$, to confirm that standard likelihood asymptotics hold for the problem at hand. This is actually the case, as shown in Figure 4, which represents the estimated distribution of the LRT, obtained with 1,000 simulated data sets under Scenario 2.

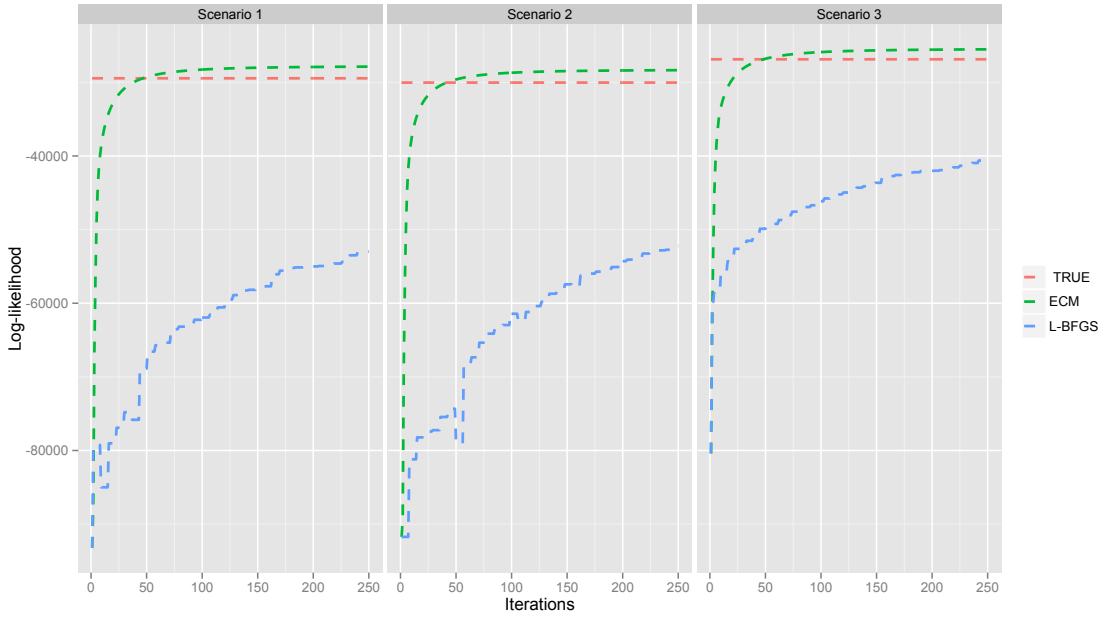

Figure 3: Log-likelihood by iteration. This graph compares the ascent of the log-likelihood function obtained with the ECM algorithm (blue line) and with the L-BFGS (green line), whereas the red line represents the log-likelihood calculated at the true parameter value. Each panel refers to a single representative data set, for each of the three scenarios.

The empirical distribution of the LRT seems close to the asymptotic distribution, and this also provides further indirect suggestion that the ECM is actually able to correctly locate the MLE.

3.2 Selection of the latent factor dimensions

Next we consider selecting the dimension of the latent space, again by means of simulation studies performed under the same three scenarios considered above. For each data set, we first choose T_s by means of standard FA techniques, and then choose K. We thus focus in particular on the problem of selecting K, tackled by applying existing model selection techniques, such as the Akaike information criterion (AIC) (Akaike, 1974) and the Bayesian information criterion (BIC) (Schwarz et al., 1978), for which there is an extensive literature (Burnham and Anderson, 2002, Preacher and Merkle, 2012). The study of model selection based on information criteria is still in progress in FA settings (Chen and Chen, 2008, Hirose and Yamamoto, 2014), so it seems useful to evaluate the behavior of both AIC and BIC for choosing K. Along the two information-based criteria, we also considered the likelihood ratio test (LRT) for choosing between nested models with different values of K.

Scenario 1 in Table 3 shows the results obtained by our model fitting approach on simulations for 100 different data sets generated independently from the MSFA with $K = 0$. The overall trend is that all the three methods favor choosing the model with $K = 0$, but both AIC and LRT do so more consistently than BIC.

Scenario 2 in Table 3 reports the results for 100 different data sets generated independently from the MSFA with $K = 1$. Here AIC outperforms both BIC and

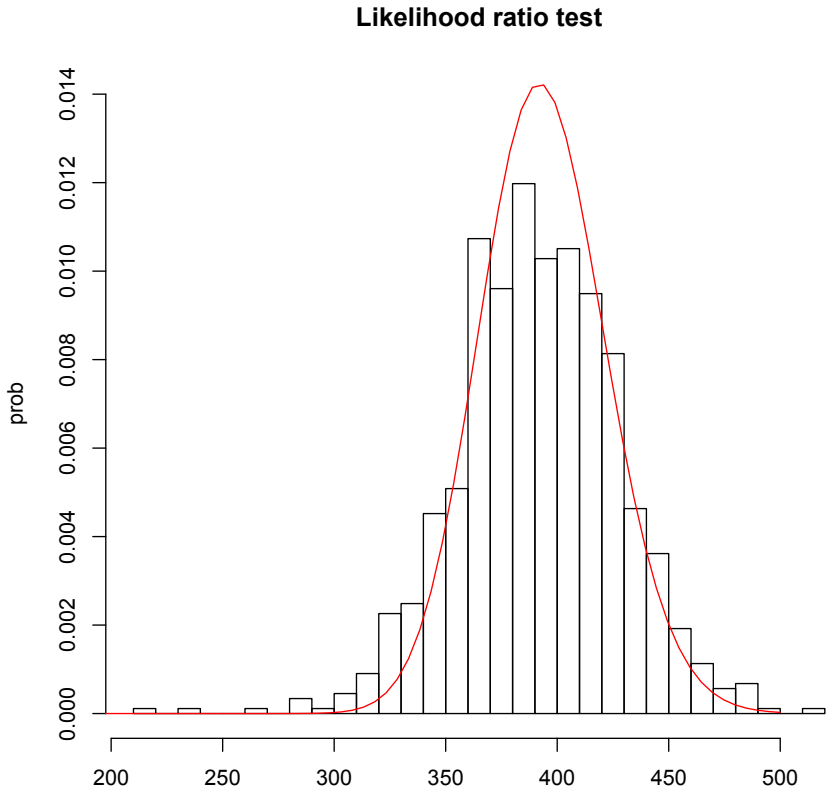

Figure 4: Distribution of the LRT for choosing between $K = 0$ and $K = 1$, obtained from 1,000 simulated data sets under Scenario 2. The red line represents the asymptotic distribution, a chi-square distribution with 192 degrees of freedom.

LRT, though the latter is not far off. The contrast with the performance of the BIC is striking. BIC never chooses a K smaller than 3 and most often chooses $K = 5$. This is a far simpler model than the data generating model. Generally, BIC penalizes model complexity more strongly than AIC, so it is not surprising that BIC tends to prefer models with more common factors and thus less parameters. In our application, subtracting a unit to K adds a large number of parameters to the model, as more factors are allowed to differ across studies.

Finally, scenario 3 in Table 3 considers the same type of comparison based on 100 different data sets are generated independently from the MSFA with $K = 3$. Again, AIC seems to be the best criterion, always leading to the selection of the true model.

The results of these three simulation studies point strongly towards the usage of AIC to select the value of K. This will therefore be the strategy employed in our applied examples.

Table 3: Comparison of model assessment methods under Scenario 1, 2 and 3.

	Method	$K=0$	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$
Scenario 1	AIC	100	0	0	0	0	0
	BIC	91	1	2	6	0	0
	LRT	100	0	0	0	0	0
	Method	$K=0$	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$
Scenario 2	AIC	0	100	0	0	0	0
	BIC	0	0	0	2	6	92
	LRT	3	97	0	0	0	0
	Method	$K=0$	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$
Scenario 3	AIC	0	0	0	100	0	0
	BIC	0	0	0	1	23	76
	Lik Ratio Test	0	0	0	91	9	0

4 Gene expression example: Immune System pathway

To illustrate MSFA in an important biological example, we analyzed the four studies described in Table 1, focusing on transcription of genes involved in immune system activity. Specifically, we considered genes included in the sub-pathways "Adaptive Immune System" (AI), "Innate Immune System" (II) and "Cytokine Signaling in Immune System" (CSI) from `reactome.org`. These sub-pathways belong to the larger Immune System pathway and do not have overlapping genes. In addition, we restricted attention to genes which are common across all studies.

Initially, we conducted preliminary analyses to asses the total latent factor dimensions, the number of common factors across studies and the number of specific factors for each study. Using the AIC, the number of common factors is set to one, and the number of specific factors for each study is as shown in Table 4.

Table 4: Number of total and specific factors for each study, Table 1, in Immune System pathway. The estimated number of common factors is 1.

Study	total latent factors	specific latent factors
GSE9891	6	5
GSE20565	7	6
GSE26712	10	9
TCGA	9	8

We then compare the cross-validation prediction errors computed by the MSFA to those computed by standard factor analysis (FA) applied separately to each study. We fit the MSFA and the standard FA on a random 80% of the data, and evaluate the prediction error on the remaining 20%. Predictions are obtained as

$$\text{MSFA:} \quad \hat{\mathbf{x}}_{is} = \hat{\mathbf{\Phi}}\hat{\mathbf{f}}_i + \hat{\mathbf{\Lambda}}_s^{\text{MSFA}}\hat{\mathbf{l}}_{is} \qquad \text{FA:} \quad \hat{\mathbf{x}}_{is} = \hat{\mathbf{\Lambda}}_s^{\text{FA}}\hat{\mathbf{l}}_{is}$$

where $\hat{\mathbf{\Lambda}}_s^{\text{MSFA}}$ are the specific factor loadings estimated with MSFA and $\hat{\mathbf{\Lambda}}_s^{\text{FA}}$ are the factor loadings estimated with FA. We evaluated the mean squared error of prediction $\text{MSE} = \frac{1}{n}\sum_{s=1}^{S}\sum_{i=1}^{n}\left(x_{is} - \hat{x}_{is}\right)^2$, using the 20% of the samples in each study set aside for each cross-validation iteration. The MSE is 10% smaller for MSFA than for FA. The MSFA is imposing a strong equality constraint on the first factor.

This analysis illustrates how MSFA borrows strength across studies in the estimation of the factor loadings, in such a way that the predictive ability in independent observation is not only preserved but even improved.

Next, we focus on the analysis of the factor themselves. The heatmap in Figure 5 depicts the estimates of the factor loadings, both the common factor (highlighted in the black rectangle) that can be identified reproducibly across the studies, and the specific ones.

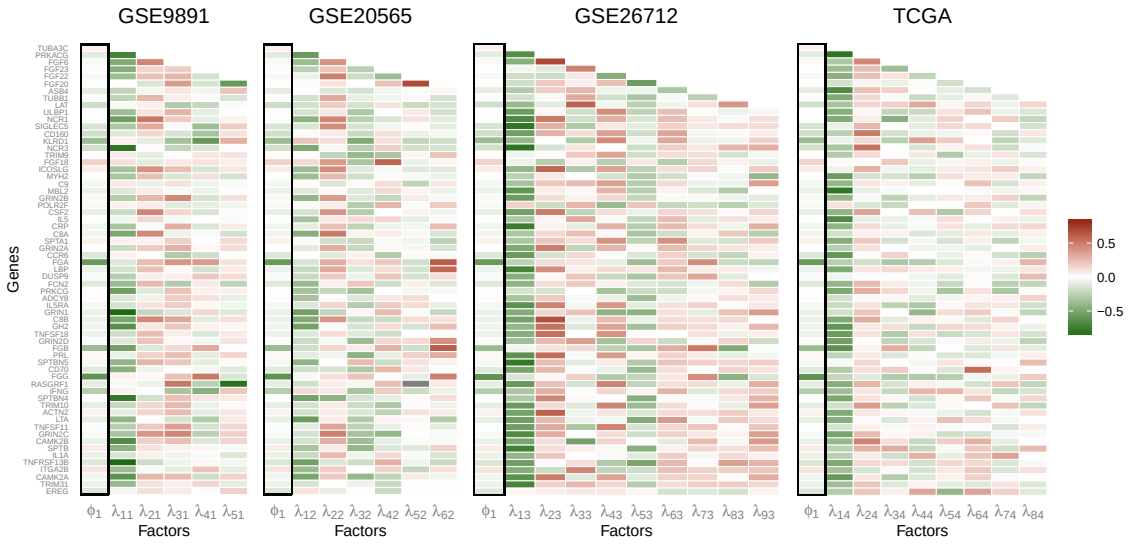

Figure 5: Heatmap of the estimated factor loadings obtained with MSFA, both common (black rectangle) and specific ones.

To help interpreting the biological meaning of the common factor, we apply Gene Set Enrichment Analysis (GSEA) for determining whether one of the three gene sets is significantly enriched among loadings that are high in absolute value (Mootha et al., 2003, Subramanian et al., 2005). We consider all the three sub-pathways in the Immune System pathway. We used the package `Rtopper` in `R` in `Bioconductor`, following the method illustrated in Tyekucheva et al. (2011).

The resulting analysis shows that the common factor is significantly enriched by the innate immune system sub-pathway, suggesting that genuine biological signal may have been identified.

Further, we consider the cross-study pairwise correlations of the loadings involving the study-specific factors shown in Figure 6. Three of the specific factors of the GSE9891 study are strongly correlated with three corresponding factors in the GSE20565 study. Absolute correlations range from 0.66 to 0.81.

To further probe this possibility, at least within the Immune System pathway, we analyze studies GSE9891 and GSE20565 separately from the other two using MSFA. The AIC chooses a model with $K = 4$ with a total of six latent factors in the former study and seven in the latter one. Studies GSE9891 and GSE20565 use the same

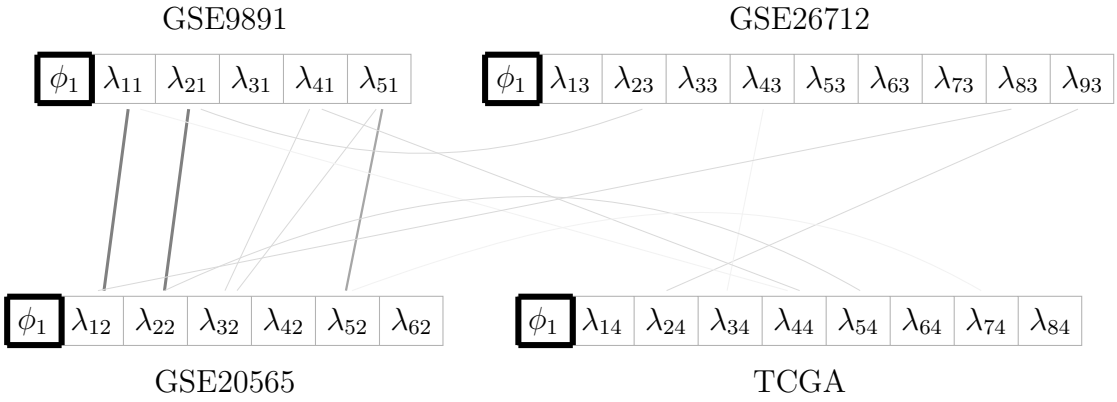
GSE9891
GSE26712
ϕ_1 λ_{11} λ_{21} λ_{31} λ_{41} λ_{51}
ϕ_1 λ_{13} λ_{23} λ_{33} λ_{43} λ_{53} λ_{63} λ_{73} λ_{83} λ_{93}
ϕ_1 λ_{12} λ_{22} λ_{32} λ_{42} λ_{52} λ_{62}
ϕ_1 λ_{14} λ_{24} λ_{34} λ_{44} λ_{54} λ_{64} λ_{74} λ_{84}
GSE20565
TCGA

Figure 6: Graphical representation of the correlation of the specific factor loadings obtained with the MSFA. Darker grey lines correspond to higher correlations. Correlations smaller than .25 are not shown.

microarray platform, Affy U133 Plus2.0, unlike the other two. This prompts the conjecture that the three stronger correlations observed may be related to technological rather than biological variation. Naturally it is also possible that there may be specific technical features of this platform that enable it to identify additional factors, although this is less likely in view of the fact that our analysis is restricted to a common set of genes.

We performed the GSEA on the estimated factor loadings for the two-study analysis. The results show that the first common factor is still related to the innate immune system pathway, as was the case for the single common factor shared between the four studies in the earlier analysis. Also, the common factor of the four studies analysis is highly correlated with the first common factors of the two-study analysis, $r = 0.60$. The three remaining common factors are not related to any of the remaining pathways, further corroborating the hypothesis that they may represent the results of spurious variation unique to the specific platform used.

We also checked the impact of the choice of a gene order, because of the dependence induced by the block lower triangular structure we assumed for $\mathbf{\Omega}_s$, to address identifiability. In particular, we repeated the same analysis was repeated after permuting. Despite minor discrepancies, the final conclusion is not changed. Namely, the one common factor is significantly enriched only with the innate immune system sub-pathway.

Overall, this analysis illustrates important features of this method, including its ability to capture biological signal common to multiple studies and technological platforms, and at the same time to isolate the source of variation coming, for example, from the different platform by which gene expression is measured.

5 Discussion

In this article we introduced and studied a novel class of factor analysis methodologies for the joint analysis of multiple studies. We hope to have provided a valuable tool to accelerate the pace at which we can combine unsupervised analysis across multiple studies, and understand the cross-study reproducibility of signal in multivariate data.

The main concept is to separately identify and estimate 1) common factors shared

across multiple studies, and 2) study-specific factors. This is intended to help address one of the most critical steps in cross-study analysis: to identify factors that are reproducible across studies and to remove idiosyncratic variation that lacks cross-study reproducibility. The method is simple and it is based on a generalized version of FA able to handle multiple studies simultaneously and to capture the two types of information.

Several methods have been proposed to analyze diverse data sets and to capture the correlation between different studies. The Co-inertia analysis (CIA) emerged in ecology to explore the common structure of two distinct sets of variables (such as species' abundances of flora and fauna) measured at the same sites (Dolédec and Chessel, 1994, Dray et al., 2003). It proceeds by separately performing dimension reduction on each set of variables, to derive factor scores for the sites. In a second, independent, stage the correlation between these factors is investigated. The Multiple Co-inertia analysis (MCIA) (Dray et al., 2003) is a generalization of CIA to consider more than two data sets. MCIA finds a hyperspace, where variables showing similar trends are projected close to each other (Meng et al., 2014).

A related method is the Multiple Factor analysis (MFA) (Abdi et al., 2013), an extension of principal component analysis (PCA) which consists of three steps. The first is a PCA for each study. In the second step each data set is normalized by dividing by its first singular value. In the third step, a single data set is created by stacking the normalized data from different studies by row, and a final PCA is done.

Two difference can be emphasized between these approaches and Multi-study Factor Analysis. First they are focused on analyzing only the common structure after having excluded the noise. Instead our method gives an estimation of both common and study-specific components. Second they operate stage-wise, decomposing each matrix separately, while our study analyzed the data jointly. This is critical in a meta-analytic context because the presence of a recognizable factor in one study can assist with the identification of the same factor in other studies even when it is more difficult to recognize it.

We develop a fast Expectation Conditional-Maximization algorithm for parameter estimates and we provide a procedure for choosing the common and specific factor. An important constraint at this stage is that the method requires more observations than there are variables. Relaxing this condition is an important direction for future work. It could be pursued in a variety of ways including regularization and Bayesian modeling.

The MSFA needs to be constrained to be identifiable and so the constraints used here is the popular block lower triangular matrix. Although this condition is largely used in classical FA settings, it induces an order dependence among the variables (Frühwirth-Schnatter and Lopes, 2010). As noted in Carvalho et al. (2008), the choice of the first $K + J_s$ variables is an important modeling decision, to be made with some care. In our application, it is somewhat reassuring that the checking made on the impact of the chosen variable order on the final conclusion leads to the same conclusions, although general conclusions cannot be drawn. Others constraints or rotation methods, such as the varimax criterion (Kaiser, 1958), could be considered, though their extension to the MSFA setting would require some further investigation.

The MSFA model can be applied to many settings when the aim is to isolate commonalities and differences across different groups, population or studies. There might be other applications where the goal is to capture study-specific features of interest and,

instead, remove common factors shared among studies. Other applications may focus on capturing both common and specific factors, without removing any of them. MSFA may have broad applicability in a wide variety of genomic platforms (e.g. microarrays, RNA-seq, SNPs, proteomics, metabolomics, epigenomics), as well as datasets in other fields of biomedical research, such as those generated by exposome studies or Electronic Medical Record (EMR). However, the concepts is simple and general enough to be of general interest across all applications of multivariate analysis.

References

Abdi, H., Williams, L. J., and Valentin, D. (2013). Multiple factor analysis: principal component analysis for multitable and multiblock data sets. *Wiley Interdisciplinary reviews: computational statistics*, 5(2):149–179.

Adachi, K. (2012). Factor analysis with EM algorithm never gives improper solutions when sample covariance and initial parameter matrices are proper. *Psychometrika*, 78(2):380–394.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Andreasen, N. C., Carpenter Jr, W. T., Kane, J. M., Lasser, R. A., Marder, S. R., and Weinberger, D. R. (2005). Remission in schizophrenia: proposed criteria and rationale for consensus. *American Journal of Psychiatry*.

Bernau, C., Riester, M., Boulesteix, A.-L., Parmigiani, G., Huttenhower, C., Waldron, L., and Trippa, L. (2014). Cross-study validation for the assessment of prediction algorithms. *Bioinformatics*, 30(12):i105–i112.

Bhattacharya, A. and Dunson, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika*, 98(2):291–306.

Blum, Y., Le Mignon, G., Lagarrigue, S., and Causeur, D. (2010). A factor model to analyze heterogeneity in gene expression. *BMC Bioinformatics*, 11(1):368.

Bonome, T., Levine, D. A., Shih, J., Randonovich, M., Pise-Masison, C. A., Bogomolniy, F., Ozbun, L., Brady, J., Barrett, J. C., Boyd, J., et al. (2008). A gene signature predicting for survival in suboptimally debulked patients with ovarian cancer. *Cancer Research*, 68(13):5478–5486.

Burnham, K. P. and Anderson, D. R. (2002). *Model Selection and Multimodel Inference: a Practical Information-theoretic Approach*. Springer, New York, second edition.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Byrne, B. M., Shavelson, R. J., and Muthén, B. (1989). Testing for the equivalence of factor covariance and mean structures: The issue of partial measurement invariance. *Psychological bulletin*, 105(3):456.

Carvalho, C. M., Chang, J., Lucas, J. E., Nevins, J. R., Wang, Q., and West, M. (2008). High-dimensional sparse factor modeling: applications in gene expression genomics. *Journal of the American Statistical Association*, 103(484):1438–1456.

Cattell, R. B. (1966). The scree test for the number of factors. *Multivariate Behavioral Research*, 1(2):245–276.

Chen, J. and Chen, Z. (2008). Extended bayesian information criteria for model selection with large model spaces. *Biometrika*, 95(3):759–771.

Cope, L., Naiman, D. Q., and Parmigiani, G. (2014). Integrative correlation: Properties and relation to canonical correlations. *Journal of Multivariate Analysis*, 123:270–280.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38.

Dolédec, S. and Chessel, D. (1994). Co-inertia analysis: an alternative method for studying species–environment relationships. *Freshwater biology*, 31(3):277–294.

Dray, S., Chessel, D., and Thioulouse, J. (2003). Co-inertia analysis and the linking of ecological data tables. *Ecology*, 84(11):3078–3089.

Edefonti, V., Hashibe, M., Ambrogi, F., Parpinel, M., Bravi, F., Talamini, R., Levi, F., Yu, G., Morgenstern, H., and Kelsey, K, e. a. (2012). Nutrient-based dietary patterns and the risk of head and neck cancer: a pooled analysis in the international head and neck cancer epidemiology consortium. *Annals of oncology*, 23(7):1869–1880.

Friguet, C., Kloareg, M., and Causeur, D. (2009). A factor model approach to multiple testing under dependence. *Journal of the American Statistical Association*, 104(488):1406–1415.

Frühwirth-Schnatter, S. and Lopes, H. F. (2010). Parsimonious Bayesian factor analysis when the number of factors is unknown. *Unpublished Working Paper, Booth Business*.

Ganzfried, B. F., Riester, M., Haibe-Kains, B., Risch, T., Tyekucheva, S., Jazic, I., Wang, X. V., Ahmadifar, M., Birrer, M. J., Parmigiani, G., et al. (2013). Curatedovariandata: clinically annotated data for the ovarian cancer transcriptome. *Database*, 2013.

Garrett-Mayer, E., Parmigiani, G., Zhong, X., Cope, L., and Gabrielson, E. (2007). Cross-study validation and combined analysis of gene expression microarray data. *Biostatistics*, 9(2):333–354.

Geweke, J. and Zhou, G. (1996). Measuring the pricing error of the arbitrage pricing theory. *Review of Financial Studies*, 9(2):557–587.

Hirose, K. and Yamamoto, M. (2014). Estimation of an oblique structure via penalized likelihood factor analysis. *Computational Statistics & Data Analysis*, 79:120–132.

Horn, J. L. (1965). A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185.

Irizarry, R. A., Hobbs, B., Collin, F., Beazer-Barclay, Y. D., Antonellis, K. J., Scherf, U., Speed, T. P., et al. (2003). Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics*, 4(2):249–264.

Jöreskog, K. G. (1967a). A general approach to confirmatory maximum likelihood factor analysis. *ETS Research Bulletin Series*, 1967(2):183–202.

Jöreskog, K. G. (1967b). Some contributions to maximum likelihood factor analysis. *Psychometrika*, 32(4):443–482.

Jöreskog, K. G. (1971). Simultaneous factor analysis in several populations. *Psychometrika*, 36(4):409–426.

Jöreskog, K. G. and Goldberger, A. S. (1972). Factor analysis by generalized least squares. *Psychometrika*, 37(3):243–260.

Kaiser, H. F. (1958). The varimax criterion for analytic rotation in factor analysis. *Psychometrika*, 23(3):187–200.

Kerr, K. F. (2007). Extended analysis of benchmark datasets for agilent two-color microarrays. *BMC Bioinformatics*, 8(1):371.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica*, 14(1):41–68.

McLachlan, G. and Krishnan, T. (2007). *The EM Algorithm and Extensions*. John Wiley & Sons, New Jersey, second edition.

Méhes, G., Luegmayr, A., Hattinger, C. M., Lörch, T., Ambros, I. M., Gadner, H., and Ambros, P. F. (2001). Automatic detection and genetic profiling of disseminated neuroblastoma cells. *Medical and Pediatric Oncology*, 36(1):205–209.

Meng, C., Kuster, B., Culhane, A. C., and Gholami, A. M. (2014). A multivariate approach to the integration of multi-omics datasets. *BMC bioinformatics*, 15(1):1.

Meng, X.-L. and Rubin, D. B. (1993). Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80(2):267–278.

Meredith, W. (1993). Measurement invariance, factor analysis and factorial invariance. *Psychometrika*, 58(4):525–543.

Meyniel, J.-P., Cottu, P. H., Decraene, C., Stern, M.-H., Couturier, J., Lebigot, I., Nicolas, A., Weber, N., Fourchotte, V., Alran, S., et al. (2010). A genomic and transcriptomic approach for a differential diagnosis between primary and secondary ovarian carcinomas in patients with a previous history of breast cancer. *BMC Cancer*, 10(1):222.

Mootha, V. K., Lindgren, C. M., Eriksson, K.-F., Subramanian, A., Sihag, S., Lehar, J., Puigserver, P., Carlsson, E., Ridderstråle, M., Laurila, E., et al. (2003). Pgc-1α-responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nature Genetics*, 34(3):267–273.

Network, C. G. A. R. et al. (2011). Integrated genomic analyses of ovarian carcinoma. *Nature*, 474(7353):609–615.

Parmigiani, G., Garrett-Mayer, E. S., Anbazhagan, R., and Gabrielson, E. (2004). A cross-study comparison of gene expression studies for the molecular classification of lung cancer. *Clinical cancer research : an official journal of the American Association for Cancer Research*, 10(9):2922–2927.

Preacher, K. J. and Merkle, E. C. (2012). The problem of model selection uncertainty in structural equation modeling. *Psychological Methods*, 17(1):1.

Riester, M., Wei, W., Waldron, L., Culhane, A. C., Trippa, L., Oliva, E., Kim, S.-H., Michor, F., Huttenhower, C., Parmigiani, G., and Birrer, M. J. (2014). Risk prediction for late-stage ovarian cancer by meta-analysis of 1525 patient samples. *Journal of the National Cancer Institute*, 106(5).

Rubin, D. B. and Thayer, D. T. (1982). EM algorithms for ML factor analysis. *Psychometrika*, 47(1):69–76.

Runcie, D. E. and Mukherjee, S. (2013). Dissecting high-dimensional phenotypes with Bayesian sparse factor analysis of genetic covariance matrices. *Genetics*, 194(3):753–767.

Scaramella, L. V., Conger, R. D., Spoth, R., and Simons, R. L. (2002). Evaluation of a social contextual model of delinquency: a cross-study replication. *Child development*, 73(1):175–195.

Scharpf, R., Tjelmeland, H., Parmigiani, G., and Nobel, A. (2009). A Bayesian model for cross-study differential gene expression. *Journal of the American Statistical Association*, 104(488):1295–1310.

Schwarz, G. et al. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.

Shi, L., Reid, L. H., Jones, W. D., Shippy, R., Warrington, J. A., Baker, S. C., Collins, P. J., De Longueville, F., Kawasaki, E. S., Lee, K. Y., et al. (2006). The microarray quality control (maqc) project shows inter-and intraplatform reproducibility of gene expression measurements. *Nature Biotechnology*, 24(9):1151–1161.

Steiger, J. H. and Lind, J. M. (1980). Statistically based tests for the number of common factors. *Paper presented at Psychometric Society Meeting, Iowa City, May.*

Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., et al. (2005). Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 102(43):15545–15550.

Thurstone, L. L. (1931). Multiple factor analysis. *Psychological Review*, 38(5):406.

Tothill, R. W., Tinker, A. V., George, J., Brown, R., Fox, S. B., Lade, S., Johnson, D. S., Trivett, M. K., Etemadmoghadam, D., Locandro, B., et al. (2008). Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome. *Clinical Cancer Research*, 14(16):5198–5208.

Tyekucheva, S., Marchionni, L., Karchin, R., and Parmigiani, G. (2011). Integrating diverse genomic data using gene sets. *Genome Biology*, 12(10):R105.

Van Loan, C. F. (2000). The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(1):85–100.

Waldron, L., Haibe-Kains, B., Culhane, A. C., Riester, M., Ding, J., Wang, X. V., Ahmadifar, M., Tyekucheva, S., Bernau, C., Risch, T., Ganzfried, B. F., Huttenhower, C., Birrer, M., and Parmigiani, G. (2014). Comparative meta-analysis of prognostic gene signatures for late-stage ovarian cancer. *Journal of the National Cancer Institute*, 106(5).

Wang, X. V., Verhaak, R., Purdom, E., Spellman, P. T., and Speed, T. P. (2011). Unifying gene expression measures from multiple platforms using factor analysis. *PloS One*, 6(3).